Bijleveld, Frits; Commandeur, Jacques; Gould, Phillip; Koopman, Siem Jan

Working Paper

Model-based Measurement of Latent Risk in Time Series with Applications

Tinbergen Institute Discussion Paper, No. 05-118/4

Provided in Cooperation with:
Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: Bijleveld, Frits; Commandeur, Jacques; Gould, Phillip; Koopman, Siem Jan (2005) : Model-based Measurement of Latent Risk in Time Series with Applications, Tinbergen Institute Discussion Paper, No. 05-118/4, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:
https://hdl.handle.net/10419/86521

TI 2005-118/4
Tinbergen Institute Discussion Paper

Model-based Measurement of Latent Risk in Time Series with Applications

Frits Bijleveld[a]
Jacques Commandeur[a]
Phillip Gould[b,c]
Siem Jan Koopman[c]

[a] *SWOV , Leidschendam, the Netherlands;*
[b] *Monash University, Melbourne, Australia;*
[c] *Vrije Universiteit Amsterdam, and Tinbergen Institute.*

Tinbergen Institute
The Tinbergen Institute is the institute for economic research of the Erasmus Universiteit Rotterdam, Universiteit van Amsterdam, and Vrije Universiteit Amsterdam.

Tinbergen Institute Amsterdam
Roetersstraat 31
1018 WB Amsterdam
The Netherlands
Tel.: +31(0)20 551 3500
Fax: +31(0)20 551 3555

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900
Fax: +31(0)10 408 9031

Please send questions and/or remarks of non-scientific nature to driessen@tinbergen.nl.
Most TI discussion papers can be downloaded at http://www.tinbergen.nl.

Model-based measurement of latent risk in time series with applications

Frits Bijleveld[(a)], Jacques Commandeur[(a)],
Phillip Gould[(b, c)] and Siem Jan Koopman[(c)*]

[(a)] *SWOV Institute for Road Safety Research, Leidschendam, Netherlands*
[(b)] *Department of Econometrics and Business Statistics, Monash University, Melbourne, Australia*
[(c)] *Department of Econometrics, Vrije Universiteit Amsterdam, Netherlands*

Abstract

Risk is at the center of many policy decisions in companies, governments and other institutions. The risk of road fatalities concerns local governments in planning countermeasures, the risk and severity of counterparty default concerns bank risk managers on a daily basis and the risk of infection has actuarial and epidemiological consequences. However, risk can not be observed directly and it usually varies over time. Measuring risk is therefore an important exercise. In this paper we introduce a general multivariate framework for the time series analysis of risk that is modelled as a latent process. The latent risk time series model extends existing approaches by the simultaneous modelling of (i) the exposure to an event, (ii) the risk of that event occurring and (iii) the severity of the event. First, we discuss existing time series approaches for the analysis of risk which have been applied to road safety, actuarial and epidemiological problems. Second, we present a general model for the analysis of risk and discuss its statistical treatment based on linear state space methods. Third, we apply the methodology to time series of insurance claims, credit card purchases and road safety. It is shown that the general methodology can be effectively used in the assessment of risk.

Keywords: Actuarial statistics; Dynamic factor analysis; Kalman filter; Maximum likelihood; Road casualties; State space model; Unobserved components.

*Corresponding author: S.J. Koopman, Department of Econometrics, Vrije Universiteit, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands. Email: `s.j.koopman@feweb.vu.nl`. This version: November 18, 2005.

1 Introduction

In the literature the term "risk" may take many meanings. In the financial econometrics literature, realized and implied volatility models treat risk as the standard deviation of returns. Governments and financial institutions are concerned with other types of risk, where particular events are studied. "Insurance risk" is modelled by actuaries who wish to measure the risk of a claim. Epidemiologists study "medical risk", which is the risk of infection or injury. For banks, "default risk" is of great importance for regulatory and internal capital management purposes. Similarly, accident researchers and credit portfolio managers are also concerned with the risk of certain events occurring, and also their likely severity. These "event risk" analyses have common elements: there is (i) exposure to risk, (ii) the risk (chance, probability) of the event occurring and (iii) the severity of the event.

The time series modelling of "event risk" offers new insights into data and can confirm or reject the validity of constant risk assumptions. There is growing pressure to develop such models in a range of fields. The Basel Accord (BIS, 2004) requires banks to be able to understand the present and also forecast the future value-at-risk of their credit portfolios. Greater regulatory capital burdens will be placed on banks who cannot demonstrate they have robust models for default risk and associated losses. Road safety researchers now have considerable pressure from governments to be able to evaluate past safety measures and forecast future accidents and injuries (WHO, 2004). The increased availability of data and continuing improvements in computer power have also opened up a range of new models which can be applied to time series data.

There is substantial evidence that simple deterministic models fail to adequately explain the dynamics of actuarial and epidemiological systems. Recently a number of articles have examined stochastically time-varying structures to model risk in epidemiological applications. For example, Dominici et al. (2004) find evidence of time varying risk factors within a generalised additive model framework used to determine the interaction between mortality rates and air pollution concentrations. The data is from 1987–1994, indicating time variation over relatively short time frames. In this approach, the natural log of fatality counts is modelled as a smooth semiparametric function of time and weather variables and a linear function of pollution levels. Finkenstadt and Grenfell (2000) find evidence of seasonal time variation in the transmission parameter for a susceptible-exposed-infected-recovered (SEIR) model for measles epidemics. An illustration of the modelling of disease incidence on the basis of unobserved or latent processes is provided by Morton and Finkenstadt (2005).

In actuarial research, there is a surprising lack of time series models for the rate and severity of insurance claims. Among the few articles is de Jong and Boyle (1983), in which Bayesian methods are applied to a state space model which produces stochastically time-varying mor-

tality rates. Exposure is measured by the size of the population and is effectively treated as a covariate. Risk is defined as the time-varying, locally polynomial probability for a binomial distribution with the number of trials given by the population size. Harvey and Fernandes (1989) also develop a model for insurance claims using the stochastic unobserved components (UC) framework (Harvey, 1989), where both the size of claims and the number of claims are modelled. The focus of the Harvey and Fernandes (1989) paper is on the modelling of Poisson and normal variables that are combined to form a loss function, which is the dollar value of all claims. Automobile insurance claims for multiple cohorts are analysed by Ledolter et al. (1990), who use stochastic UC models to test for common factors ("shrinkage") across cohorts. These models provide better fit to the data than static or deterministic models.

In bank risk management there have also been a number of articles examining the use of time varying parameters to model the risk of counterparty default. Allen and Saunders (2003) highlight the need for dynamic approaches to modelling company default. Structural financial risk models have been used successfully in this area. Multivariate approaches are often used to model observations from different cohorts, rather than separate "dimensions" of risk such as exposure to default and the number of defaults. A time-varying logistic model estimated via Kalman filtering is introduced by Fahrmeir and Wagenpfeil (1996) for duration models which assess the probability of subjects entering or leaving a state of unemployment. In this case the model is for individuals, and time is the only measure of exposure, so a logistic transform is used to produce time varying probabilities. The results suggest there is a need for time variation in model parameters to accurately model unemployment dynamics.

In road safety research, cross-sectional induced exposure methods (Li and Kim, 2000) have been used to separate the effects of crash risk and exposure in the absence of exposure data. However, time series methods such as the *demande routière, des accidents et leur gravité* (DRAG) framework of Gaudry (1984) and Gaudry and Lassarre (2000) has been applied more widely. The DRAG framework aims at modelling the dimensions of exposure, risk and severity. As in most "event risk" models, exposure is a covariate used to explain accident counts. The DRAG model differs from other structures in that it is usually estimated using ARMA regression models with many social and economic explanatory variables. The DRAG model's reliance on economic and social variables, such as unemployment and alcohol consumption (which are difficult to forecast in their own right), means that it is of limited value in a forecasting context. The UC time series framework is adopted in the road safety study of Harvey and Durbin (1986).

Despite the many common features of the data and methods applied, to date there is no single unified framework for modelling "event risk". In this paper we introduce a general multivariate model for "event risk" analysis that can consider exposure, risk and severity simultaneously. The latent risk time series (LRT) model can be applied to a range of problems involving "event risk" and is not specifically limited to either actuarial or epidemiological ap-

plications. The standard approaches to risk treat exposure as a fixed and known variable and few models allow for the modelling of severity at all. The LRT model is general and allows for the stochastic evolution of exposure, risk and severity over time. It extends previous work by treating exposure and severity as an integral part of the risk problem. The logarithmic transformation is used to model multiplicative, time-varying relationships between exposure, risk and severity. In existing approaches some or all of these variables (particularly exposure) are treated as known, when in reality they are often measured under error and subject to stochastic variation. The LRT model has a multivariate structure and therefore correlations between latent processes and errors can be estimated. Since the model belongs to the class of UC time series models, the multivariate decomposition can include latent factors for trend, seasonal and cyclical dynamics. This general framework also allows for the forecasting of future exposures, events and losses together with prediction confidence bounds, which are of particular interest to risk managers. The relative simplicity of the model structure means that complex estimation techniques are usually not required. Bayesian and classical estimation methods can be easily applied and usually rely on Kalman filter methods. The multivariate framework is sufficiently general to allow for a study of multiple cohorts with a view to enabling "shrinkage" of the number of states required, see also Ledolter et al. (1990).

The statistical framework, including state space forms and estimation methods are presented in Section 2. The exposure-risk motor vehicle insurance model is the first example of a LRT analysis and is discussed in Section 3. The exposure-risk-severity model for credit card use is treated in Section 4. The multiple exposure-single risk model for bicycle and moped road traffic accidents is presented in Section 5. The LRT analysis includes parameter estimation, measurement of exposure, risk and severity (where applicable) based on signal extraction methods and forecasts. Section 6 concludes.

2 The statistical framework

The latent risk time series (LRT) model includes latent factors for exposure, risk and severity which are all associated with the observed variables:

- x_{it}, exposure at time t for group i with $i = 1, ..., k_x$
- y_{it}, outcome at time t for group i with $i = 1, ..., k_y$
- z_{it}, loss at time t for group i with $i = 1, ..., k_z$,

for $t = 1, \ldots, n$ where n is the number of observations and where k_x, k_y and k_z are the number of groups for exposure, outcome and loss variables, respectively. There is no need to set

$k_x = k_y = k_z$ because multiple outcomes for only one exposure variable (such as multiple types of vehicle accidents per registered vehicle) can exist so that $k_x < k_y$.

The exposure variable may take many forms. For example, exposure to traffic crash risk may be measured in terms of vehicle registrations or by distance travelled. For a bank, the exposure variable may be either the number of loans or the dollar value of loans. The exposure is to the risk of a particular event or outcome occurring. The outcome variable y_{it} is typically the number of times a certain event occurs within group i. In insurance applications this may be the number of claims. For epidemiologists the outcome variable may be the number of successful treatments. Indeed, the outcome variable does not have to be undesirable.

In some cases a loss variable is also required to fully describe the risk associated with a particular outcome. The loss variable measures the severity (or consequences) of the outcomes. For an insurance company the loss variable may be the dollar value of claims. The severity of a loss is also important for bank risk managers who wish to know how much money is likely to be lost in the event of default.

2.1 The LRT model

The multivariate unobserved components time series modelling framework is adopted to formulate a risk system for the observed variables exposure, outcome and loss. The latent risk model (LRT) model relates these observed variables within a multivariate system of equations:

$$\begin{aligned} x_{it} &= E_{it} \times U_{it}^{(x)}, \\ y_{it} &= E_{it} \times R_{it} \times U_{it}^{(y)}, \\ z_{it} &= E_{it} \times R_{it} \times S_{it} \times U_{it}^{(z)}, \end{aligned}$$

where E_{it}, R_{it}, and S_{it} are the latent variables exposure, risk and severity for group i at time t, respectively. A multiplicative system is presented that has the expected outcome as a proportion (risk) of exposure while loss is a multiple (severity) of the expected outcome. The multiplicative error terms $U_{it}^{(x)}$, $U_{it}^{(y)}$ and $U_{it}^{(z)}$ reflect that observed variables are measured under uncertainty. For example, road-use surveys may not be accurately reported or there may simply be unexplainable short-term deviations from the trend.

After taking logs, the multiplicative LRT equations can be expressed in additive form:

$$\begin{aligned} \log x_{it} &= \mu_{it}^{(E)} + \varepsilon_{it}^{(x)}, \\ \log y_{it} &= \mu_{it}^{(E)} + \mu_{it}^{(R)} + \varepsilon_{it}^{(y)}, \\ \log z_{it} &= \mu_{it}^{(E)} + \mu_{it}^{(R)} + \mu_{it}^{(S)} + \varepsilon_{it}^{(z)}, \end{aligned} \tag{1}$$

where $\mu_{it}^{(E)} = \log E_{it}$ is log-exposure, $\mu_{it}^{(R)} = \log R_{it}$ is log-risk and $\mu_{it}^{(S)} = \log S_{it}$ is log-severity for group i and time t. The additive noise terms $\varepsilon_{it}^{(x)} = \log U_{it}^{(x)}$, $\varepsilon_{it}^{(y)} = \log U_{it}^{(y)}$ and $\varepsilon_{it}^{(z)} = \log U_{it}^{(z)}$

are identically independently distributed (i.i.d.) with mean zero but their distributions do not necessarily need to be determined. The latent factors of exposure, risk and severity have time subscripts and the form of time-variation is established below.

Most risk models implicitly assume that the relation between outcome and exposure is time-invariant. For example, the ratio of outcome and exposure, y_{it}/x_{it} or its log-ratio counterpart $\log y_{it} - \log x_{it}$, is often directly modelled as the observed series. Alternatively, in the volume edited by Gaudry and Lassarre (2000) (a road accident example) and Ledolter et al. (1990) (an insurance example), the authors use exposure as an explanatory variable and they find a negative causal relationship between outcome and exposure. The LRT model treat changes in outcome as the result of changes in the latent processes of both exposure and risk. Observation noise is treated separately for each equation and the dynamics of the different factors are modelled explicitly. The attractive feature of the LRT model is that exposure, risk and severity factors are treated simultaneously in this multivariate framework where factors are time-varying and can have common dynamics.

Unobserved components (UC) models provide a flexible framework for allowing both deterministic and stochastic variation in the latent factors. A possible UC model to describe the evolution of the latent factors is the multivariate local linear trend (LLT) specification, see Harvey (1989). The LLT model for the log-exposure factor $\mu_t^{(E)}$ is given by

$$\begin{array}{rclcrcl} \mu_t^{(E)} & = & \mu_{t-1}^{(E)} & + & \delta_{t-1}^{(E)} & + & \eta_t^{(E)}, \quad \eta_t^{(E)} \sim \text{i.i.d.}(0, \sigma^2_{\eta^{(E)}}), \\ \delta_t^{(E)} & = & \delta_{t-1}^{(E)} & & & + & \zeta_t^{(E)}, \quad \zeta_t^{(E)} \sim \text{i.i.d.}(0, \sigma^2_{\zeta^{(E)}}), \end{array} \tag{2}$$

for $t = 1, \ldots, n$ and with the subscript i suppressed for convenience of notation. When exposure consists of different groups, the subscript i can be reintroduced. The trend component $\mu_t^{(E)}$ is a random walk process with stochastic drift $\delta_t^{(E)}$ and trend innovation $\eta_t^{(E)}$ where the latter is i.i.d. with mean zero and variance $\sigma^2_{\eta^{(E)}}$. The drift term is modelled as a random walk process with i.i.d. innovation $\zeta_t^{(E)}$ that has mean zero and variance $\sigma^2_{\zeta^{(E)}}$. The trend and slope innovations are mutually uncorrelated at all time points. Some special cases of the LLT model are the linear trend function ($\sigma^2_{\eta^{(E)}} = \sigma^2_{\zeta^{(E)}} = 0$), the random walk ($\delta_t^{(E)} = 0$ for all t), the random walk with fixed drift ($\sigma^2_{\zeta^{(E)}} = 0$) and the integrated random walk or smooth trend ($\sigma^2_{\eta^{(E)}} = 0$). An important assumption is that all additive errors in the equations are serially uncorrelated. The LLT model can also be formulated for log-risk $\mu_t^{(R)}$ and log-severity $\mu_t^{(S)}$. Further, when different factors exist for different groups, the trend and drift components are formulated for different groups and their corresponding errors can be correlated contemporaneously.

In case the observations are subject to seasonal and cyclical fluctuations, associating seasonal and cyclical components need to be included in the observation equations of the LRT model (1). The seasonal component can be modelled by a series of trigonometric functions while the cyclical component can be specified as an autoregressive process with complex roots in

its polynomial, see Harvey (1989, §2.3.4) for more details on the dynamic specification of unobserved components. The seasonal and cyclical processes vary stochastically over time and are usually assumed to be driven by i.i.d. innovations with mean zero and a finite variance. Furthermore, regression effects (based on covariates or explanatory variables) can be included in the model. Such unknown deterministic effects can be added to the model when significant deviations from the exposure-risk-severity relationships can be captured by specific covariates.

For a bivariate exposure-risk model, seasonal variations in road use often appear as persistent periodic variations in exposure and risk. In this case, seasonal components for exposure and risk can be added to the measurement equations for exposure and outcome, that is

$$\begin{array}{rcccccccc} \log x_t &=& \mu_t^{(E)} &+& \gamma_t^{(E)} &+& & & \varepsilon_t^{(x)}, \\ \log y_t &=& \mu_t^{(E)} &+& \gamma_t^{(E)} &+& \mu_t^{(R)} &+& \gamma_t^{(R)} &+& \varepsilon_t^{(y)}, \end{array} \tag{3}$$

where $\gamma_t^{(E)}$ and $\gamma_t^{(R)}$ are the seasonal components for exposure and risk, respectively, In this case, the seasonal factors are an integral part of the exposure-risk system of equations. Alternatively, the seasonal components can be exclusively attributed to a particular observation equation. A seasonal component is then added exclusively to an observation equation while the other observation equations are not affected by this component. The alternative formulation with seasonal components is then given by

$$\begin{array}{rcccccccc} \log x_t &=& \mu_t^{(E)} &+& & & \gamma_t^{(x)} &+& \varepsilon_t^{(x)}, \\ \log y_t &=& \mu_t^{(E)} &+& \mu_t^{(R)} &+& \gamma_t^{(y)} &+& \varepsilon_t^{(y)}. \end{array} \tag{4}$$

where $\gamma_t^{(x)}$ and $\gamma_t^{(y)}$ represent seasonal time series processes. The seasonal variation in $\log y_t$ is captured by $\gamma_t^{(y)}$ rather than $\gamma_t^{(E)} + \gamma_t^{(R)}$. The innovations of the seasonal processes $\gamma_t^{(x)}$ and $\gamma_t^{(y)}$ may still be correlated contemporaneously. As a consequence, only $\mu_{it}^{(E)}$ represents log-exposure in the observation equation for $\log y_t$. The same discussion of alternative LRT model specifications applies to cyclical components and regression effects.

A final concern is the incorporation of intervention effects in the LRT model. Intervention variables are used to model outlying observations and breaks in the trend and drift components. An illustration of intervention analysis in the context of unobserved components time series models is given by Harvey and Durbin (1986). They investigate the effect of seat belt legislation in the UK on road casualties. Interventions can be incorporated as follows. We adopt this approach to interventions and apply it to a multivariate LRT model. An outlier at time s in the outcome variable y can be captured by considering the observation equation

$$\log y_t \;=\; \mu_t^{(E)} \;+\; \mu_t^{(R)} \;+\; D_t^O(s)\beta_s^{(y,O)} \;+\; \varepsilon_t^{(y)}, \qquad t = 1, \ldots, n, \tag{5}$$

where $D_t^O(s)$ takes the zero value at all time points except at $t = s$ where it is unity while $\beta_s^{(y,O)}$ measures the effect of this outlier intervention on the outcome variable y. Similar outlier

effects can be considered for other time points and for the other observation equations. The intervention variable $D_t^L(s)$ is for a level change or break and takes the zero value for $t = 1, \ldots, s-1$ while it is unity for $t = s, \ldots, n$. The intervention variable $D_t^D(s)$ is for a drift change or break and is zero for $t = 1, \ldots, s-1$ while it takes the value $t-s+1$ for $t = s, \ldots, n$. Their effects are measured by $\beta_s^{(y,L)}$ and $\beta_s^{(y,D)}$ for level and drift breaks, respectively. In this specification, it is assumed that the interventions only affect a particular observed variable x, y or z. Such an approach is useful when an intervention has changed the way the data is observed rather than affecting the unobserved components. Alternatively, the level and drift interventions may instead affect a particular unobserved factor (exposure, severity or risk). For example, in the case of the trend component of exposure, the level and drift interventions can be incorporated by

$$\begin{array}{rclcrcrcl} \mu_t^{(E)} & = & \mu_{t-1}^{(E)} & + & \delta_{t-1}^{(E)} & + & D_t^O(s)\beta_s^{(E,L)} & + & \eta_t^{(E)}, \quad \eta_t^{(E)} \sim \text{i.i.d.}(0, \sigma^2_{\eta^{(E)}}), \\ \delta_t^{(E)} & = & & & \delta_{t-1}^{(E)} & + & D_t^O(r)\beta_r^{(E,D)} & + & \zeta_t^{(E)}, \quad \zeta_t^{(E)} \sim \text{i.i.d.}(0, \sigma^2_{\zeta^{(E)}}), \end{array} \tag{6}$$

for different time points s and r, where $D_t^O(s)$ is defined above and $\beta_s^{(E,L)}$ and $\beta_s^{(E,D)}$ measure the effect of level and drift interventions for exposure at time s, respectively. Note that the outlier dummy variable $D_t^O(s)$ is correctly used since the level and drift equations are recursive so that a single non-zero value in $D_t^O(s)$ has a permanent effect on the level and drift components. The model specification determines whether the intervention has an effect on the observation variable directly or on one of the unobserved trend factors. The appropriate model specification will vary from case to case and is typically the result of a modelling process. The estimation and testing of intervention effects are carried out in the same way as for regression effects.

In case exposure, outcome and loss are observed for different groups, the LRT model (1) is extended to have the unobserved component of severity. When data exists for multiple groups correlated unobserved components for different groups can be analysed. Different poolings within groups can take place in a straightforward way. The applications in this paper will show that the LRT framework is flexible. Further, correlations between the exposure, risk and severity components can also be specified to examine relationships between components or to detect and to test for common factors. In particular, for an exposure-risk-severity model we can define

$$\mu_t = \begin{pmatrix} \mu_t^{(E)} \\ \mu_t^{(R)} \\ \mu_t^{(S)} \end{pmatrix}, \quad \delta_t = \begin{pmatrix} \delta_t^{(E)} \\ \delta_t^{(R)} \\ \delta_t^{(S)} \end{pmatrix}, \quad \eta_t = \begin{pmatrix} \eta_t^{(E)} \\ \eta_t^{(R)} \\ \eta_t^{(S)} \end{pmatrix}, \quad \zeta_t = \begin{pmatrix} \zeta_t^{(E)} \\ \zeta_t^{(R)} \\ \zeta_t^{(S)} \end{pmatrix}, \quad \varepsilon_t = \begin{pmatrix} \varepsilon_t^{(x)} \\ \varepsilon_t^{(y)} \\ \varepsilon_t^{(z)} \end{pmatrix}, \tag{7}$$

for $t = 1, \ldots, n$. The LRT model (1) can be expressed as

$$\begin{pmatrix} \log x_t \\ \log y_t \\ \log z_t \end{pmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \mu_t + \varepsilon_t, \qquad \mu_t = \mu_{t-1} + \delta_{t-1} + \eta_t, \qquad \delta_t = \delta_{t-1} + \zeta_t, \tag{8}$$

for $t = 1, \ldots, n$. Correlations within the individual elements of the vector components are introduced by considering

$$\eta_t \sim \text{i.i.d.}(0, \Sigma_\eta), \qquad \zeta_t \sim \text{i.i.d.}(0, \Sigma_\zeta), \qquad \varepsilon_t \sim \text{i.i.d.}(0, \Sigma_\varepsilon), \tag{9}$$

for $t = 1, \ldots, n$. The variance matrices Σ_η (level noise), Σ_ζ (slope noise) and Σ_ε (observation noise) can be non-diagonal so that correlations can exist within the latent factors μ_t (level), δ_t (slope) and ε_t, respectively, but not between them. These correlations can produce more accurate prediction intervals for exposure, outcome and loss variables and this is useful for risk managers. In the case the LRT model also includes seasonal and cyclical components, a similar specification applies to the innovations that drive the dynamics of these components. The innovations are allowed to be mutually and contemporaneously correlated. For example, seasonal variations in exposure and risk, or in purchases and expenditures, may be highly correlated. In this case, denote the seasonal innovation by ω_t where

$$\omega_t = \begin{pmatrix} \omega_t^{(E)} \\ \omega_t^{(R)} \\ \omega_t^{(S)} \end{pmatrix}, \quad \text{in case of (3)}, \qquad \omega_t = \begin{pmatrix} \omega_t^{(x)} \\ \omega_t^{(y)} \\ \omega_t^{(z)} \end{pmatrix}, \quad \text{in case of (4)},$$

that is distributed as

$$\omega_t \sim \text{i.i.d.}(0, \Sigma_\omega), \qquad t = 1, \ldots, n,$$

where Σ_ω is the variance matrix that is typically non-diagonal.

2.2 State space analysis

The unobserved components equations of the LRT model can be written in state space form. The general state space form for a vector of observed variables is given by

$$\mathbf{y_t} = \mathbf{Z_t}\alpha_\mathbf{t} + \mathbf{e_t}, \qquad \mathbf{e_t} \sim \text{i.i.d.}(0, \mathbf{H_t}), \tag{10}$$

$$\alpha_{\mathbf{t+1}} = \mathbf{T_t}\alpha_\mathbf{t} + \mathbf{u_t}, \qquad \mathbf{u_t} \sim \text{i.i.d.}(0, \mathbf{Q_t}), \tag{11}$$

where $\alpha_\mathbf{t}$ is the state vector, $\mathbf{e_t}$ is the disturbance vector of the observation equation (10) and $\mathbf{u_t}$ is the disturbance vector of the state equation (11). The disturbance vectors are mutually and serially uncorrelated at all time points. The system matrices of the state space form are

$\mathbf{Z_t}$, transition matrix $\mathbf{T_t}$ and the variance matrices $\mathbf{H_t}$ and $\mathbf{Q_t}$. The system matrices can be time-varying but are usually time invariant and sparse matrices. Some elements of the system matrices are unknown and are treated as parameters that can be estimated by maximum likelihood.

As an illustration we show how the LRT model (8) can be represented in the state space form (10) – (11). The LRT model has $k = 1$ (the subscript i is suppressed), unobserved factors for exposure, risk and severity and no seasonal component. The observation vector $\mathbf{y_t}$, the state vector $\alpha_\mathbf{t}$ and disturbance vectors for the LRT model are given by

$$\mathbf{y_t} = \begin{pmatrix} \log x_t \\ \log y_t \\ \log z_t \end{pmatrix}, \qquad \alpha_\mathbf{t} = \begin{pmatrix} \mu_t \\ \delta_t \end{pmatrix}, \qquad \mathbf{e_t} = \varepsilon_t, \qquad \mathbf{u_t} = \begin{pmatrix} \eta_t \\ \zeta_t \end{pmatrix},$$

respectively, where the vector elements of $\alpha_\mathbf{t}$ and $\mathbf{u_t}$ are given in (7). The state space form of the LRT model has the system matrices

$$\mathbf{Z_t} = (1\ 0) \otimes \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 1 & 1 \end{bmatrix}, \quad \mathbf{T_t} = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix} \otimes I_3, \quad \mathbf{H_t} = \Sigma_\varepsilon, \quad \mathbf{Q_t} = \begin{bmatrix} \Sigma_\eta & 0 \\ 0 & \Sigma_\zeta \end{bmatrix},$$

where the block elements of $\mathbf{H_t}$ and $\mathbf{Q_t}$ correspond to the variance matrices in (9). In this representation the state vector $\alpha_\mathbf{t}$ contains the latent processes for exposure, risk and severity. In more general settings, the state vector also contains seasonal, cyclical and regression components when they are included. Matrix $\mathbf{Z}$ selects the latent processes from the state vector that relate to the observation variables. Also it determines whether a component takes part of the exposure-risk-severity system or whether it is used to capture deviations from the system (see the earlier discussions). The transition matrix $\mathbf{T}$ determines the dynamic properties of the latent factors. The variance matrix $\mathbf{Q_t}$ is block diagonal. When the state vector contains other components, $\mathbf{Q_t}$ remains block diagonal so that there is no correlation between level, drift and seasonal components. Correlations only exist between state elements of the same "type" (e.g. between the drift terms of different groups).

2.3 Measurement and estimation

The state space framework contains two unknown entities, the state vector $\alpha_\mathbf{t}$ and the parameter vector ψ that contains all unknown elements in the system matrices. The state vector can be predicted conditional on observations using so-called state space methods. Filtering refers to the estimation of $\alpha_\mathbf{t}$ conditional on $\mathbf{y_1}, \ldots, \mathbf{y_t}$, that is all observations up to and including $\mathbf{y_t}$. Smoothing is similar but the estimation is conditional on all observations $\mathbf{y_1}, \ldots, \mathbf{y_n}$. Filtering and smoothing methods also compute standard errors of the estimates. The Kalman filter and

related methods carry out the necessary computations for the linear state space model (10) – (11). In case the disturbances are normally distributed, we obtain minimum mean squared estimators. When normality is not assumed, they are minimum mean squared linear estimators. A more detailed discussion on these matters can be found in Durbin and Koopman (2001).

The Kalman filter carries out the prediction error decomposition for a given state space model (10) – (11) and a particular value of ψ. This implies that the likelihood function can be evaluated by the Kalman filter for a given ψ. Maximum likelihood estimation of ψ is therefore a standard exercise of numerically maximising the likelihood function with respect to ψ and for which the Kalman filter is used to evaluate the likelihood function for different ψ's in a computationally efficient way, see Harvey (1989). In the case of the LRT model (8), parameter estimation is limited to the variance matrices $\mathbf{\Sigma}_\varepsilon$, $\mathbf{\Sigma}_\eta$ and $\mathbf{\Sigma}_\zeta$. To ensure positive semi-definite variance matrices, a particular variance matrix $\mathbf{\Sigma}_\cdot$ is decomposed by $\mathbf{\Sigma}_\cdot = M'M$ where M is a symmetric matrix. Estimation concentrates on M. Monte Carlo methods are employed to produce confidence intervals for the parameter estimates. For this purpose, matrices M are simulated from the multivariate normal distribution with mean vech($\widehat{M}$) and variance matrix V where $\widehat{M}$ is the maximum likelihood estimate of M and V equals the negative inverse of the numerical second derivative of the likelihood function with respect to vech(M) and evaluated at $M = \widehat{M}$. Elements of M that tend to get very close to zero during the estimation process are fixed at zero when some benchmark is reached.

3 Case I: a two-dimensional insurance LRT model

A two-dimensional LRT model is considered for the measurement of exposure and risk in relation to motor vehicle fatality insurance claims. An annual dataset from Victoria, Australia is analysed for this purpose. Victoria vehicle registrations represent the number of policies (exposure) for the compulsory third-party insurance body, the Transport Accident Commission (TAC). The fatality series is the outcome variable and represents the numbers of claims on these policies for different years. Although the TAC was created in 1986, the data used to fit the model is from 1950–2001. The earlier observations are valid for parameter estimation as they represent the same portfolio, which is the state's entire vehicle population. The LRT model disentangles exposure and risk effects from observations of vehicle registrations and of motor vehicle fatalities. This type of data has previously been analysed using the Oppe model (Oppe, 1989), which assumes exposure follows a logistic-S curve and log-risk evolves deterministically and typically with a downwards drift.

The observed series are presented in Figure 1. The exposure series, registrations, displays an upwards trend and is generally smooth. This is typical of the evolution of vehicle registrations in developed economies. The fatal claims series has a "hump" shape, with a peak during the

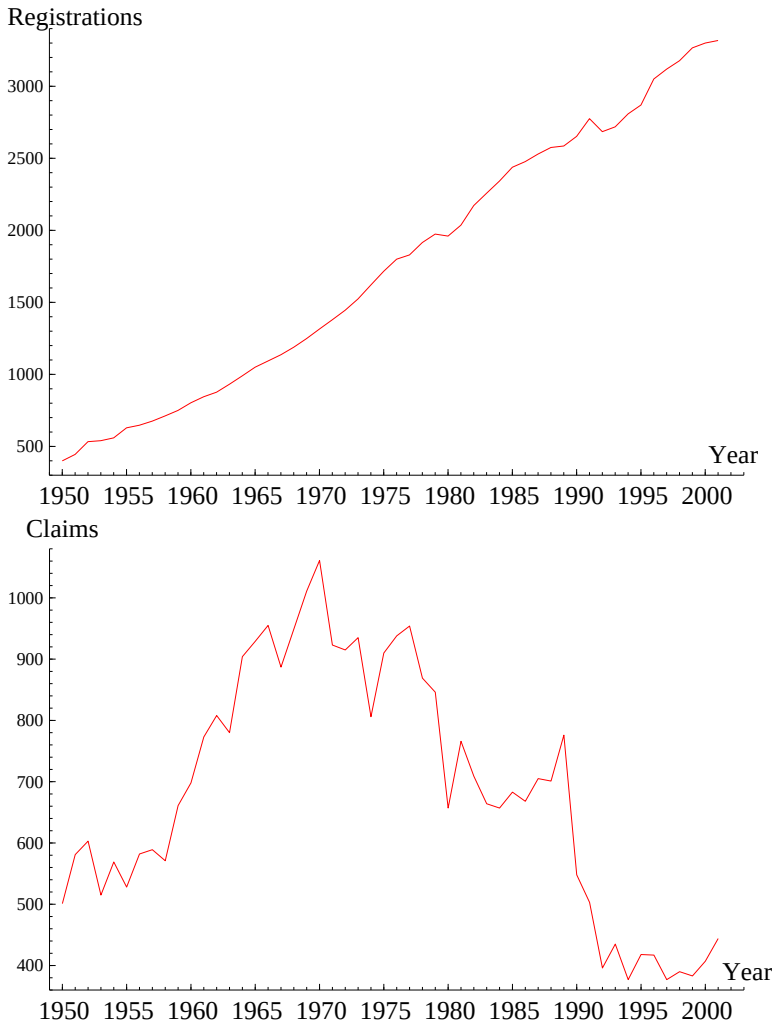

Figure 1: Victorian registrations and vehicle crash fatalities 1950–2001.

start of the 1970s, which is also typical of developed economies. Because registrations have increased monotonically over the past 50 years, the reduction in fatal claims must have been caused by a decrease in risk. Risk reductions have been driven by gradual improvements in vehicle and road design together with increased public awareness. Demographic factors may also be important as a new generation of road users ("baby boomers") began to start driving. Public horror at a road toll of 1034 in 1970 led to newspaper declarations of "war on 1034", indicative of these changing attitudes towards road safety. The effects on attitude have proved to be long-term and not only instantaneous. Other important relevant events in the sample are the introduction of seat belt laws in 1971 and the increased enforcement of mass media advertising campaigns on road safety and public safety consciousness in the early 1990s.

Because the data is not separated into groups or cohorts, there are only two observed series: policies x_t and claims y_t (the group index i is suppressed from the notation). In many insurance

examples, separate cohorts for age or gender are used to model risk and exposure for different policy holders. If cohorts were used to separate male and female policy holders, we would require four observed series (two for exposure and two for fatalities) to produce estimates of two exposure and two risk processes. Furthermore, a three-dimensional LRT can be considered to model the severity of the claims. In this case a third observed variable (the dollar value of payouts on claims) is needed to explain the expected cost per claim (severity).

In this section we focus on the two-dimensional and single cohort structure described by

$$\begin{array}{rcccl} \text{policies}_t & = & x_t & = & E_t \times U_t^{(x)} \\ \text{claims}_t & = & y_t & = & E_t \times R_t \times U_t^{(y)}, \end{array} \tag{12}$$

where E_t is exposure, R_t is risk and $U_t^{(x)}$ and $U_t^{(y)}$ are observation noise terms for policies and claims, respectively. The observations are subject to the log-transformation and the unobserved components $\mu_t^{(E)} = \log E_t$ and $\mu_t^{(R)} = \log R_t$ are flexibly modelled by the LLT specification (2). Covariates in the form of dummy variables for special events are introduced in the level ($\mu_t^{(E)}$ and $\mu_t^{(R)}$) and drift ($\delta_t^{(E)}$ and $\delta_t^{(R)}$) equations. The following events are considered: (i) in 1970, publicity started to have a safer attitude in traffic ("war on 1034"); (ii) in 1971, introduction of seat belt laws; (iii) in 1980, change in data collection on vehicle registrations; (iv) in 1990, enforcement of advertising initiatives; (v) in 1992, another change in data collection on vehicle registrations. The changes in data collection should only affect exposure E_t while the other events should have an effect on risk R_t. The change of attitude in traffic is a long-term effect and therefore captured by a change in the drift term of risk. The effects of seat belt laws and intensified road safety advertisements are taken as immediate step changes in the level of risk. Other interventions can also be considered but they have proved to be less important in the analysis.

Parameter estimates and simulated asymmetric 95% confidence intervals are presented in Table 1. The estimate of the variance of the observation noise for the registrations equation (policies) is nearer to zero than for the claims equation. This can be explained by the fact that data collection of registrations is done more accurately and the time series of claims is subject to more observation noise, see Figure 1. An interesting result from this analysis is that the level processes for exposure and risk are perfectly negatively correlated and that the drift processes are also perfectly negatively correlated. In other words, the latent variables exposure and risk are driven by two univariate noise sequences for level η_t^c and drift ζ_t^c. It follows that

$$\begin{array}{rcccccl} \mu_t^{(E)} & = & \mu_{t-1}^{(E)} & + & \delta_t^c & + & \eta_t^c, \qquad \delta_t^c = \delta_{t-1}^c + \zeta_t^c, \\ \mu_t^{(R)} & = & \mu_{t-1}^{(R)} & + & a\delta_t^c & + & b\eta_t^c, \end{array}$$

for coefficients $a < 0$ and $b < 0$ with common drift term δ_t^c and common noise terms η_t^c and ζ_t^c for level and drift, respectively. The perfect negative correlations mean that both exposure and

risk components are subject to the same stochastic shocks that determine their time-varying behaviour. This finding is in agreement with most road crash research which finds a strong negative relationship between risk and exposure. The perfect correlation of shocks implies that the components can be interpreted as common factors. Nevertheless, the estimated components are distinct from each other since they are also subject to a number of interventions that are captured by dummy variables associated with special events. The interventions allow the components to have separate shocks or breaks despite the fact that the random shocks are perfectly correlated. The regression estimates of these intervention coefficients are summarised in Table 2. The estimated intervention for the anticipated break in the level of exposure due to a change in the data collection of policies (registrations) is clearly significant for 1992 but less significant for 1980. The level interventions for risk in 1971 (seat belt laws) and 1990 (advertising initiatives) are very significant. Finally, the drift intervention for risk in 1970 (change in public safety consciousness initiated by government and media) is also very significant. All estimated interventions have negative values.

Figure 2 presents the estimated level and drift components of exposure and risk (in logs). The estimated components are subject to both random shock and interventions. The salient features of the analysis are the increasing exposure with a significant drift term throughout the sample, and the decreasing risk with a significant negative drift term that is mainly caused by the publicity intervention. Risk displays relatively more stochastic variation in both the estimated level and drift terms. The level also has a relatively large variance estimate of 0.00130. It is interesting to detect in Figure 2 that, apart from the intervention shocks, level and drift components of risk are perfectly and negatively correlated with level and slope components of exposure, respectively. For example, it means that as the positive slope of exposure becomes less positive, the negative slope for risk also becomes less negative. It follows that the slopes of risk and exposure are of opposite sign but both evolve closer towards zero. This suggests a long-term flattening of risk and exposure which is evident in the data. The level terms are also perfectly and negatively correlated. As exposure increases around its slope, risk decreases. Exposure evolves relatively smoothly, with the slope term driving much of the variation.

The slope of risk becomes significantly negative after interventions in the early 1970s. The advertising and enforcement initiatives of the early 1990s were also highly effective in reducing risk. Major variations in risk are driven by interventions together with some relatively smooth stochastic variation. Since the early 1990s, log-risk has settled to a relatively steady level, with some evidence of a recent increase. Confidence intervals for risk are wider than for exposure. It indicates that most of the variation in claims is risk-driven. The intervention breaks affect periods where there is a breakdown in the perfect correlation. A clear example is the estimated drift terms that are perfectly negatively correlated. The seat belt law shift of 1971 in the drift equation of risk means that the estimated slope terms appear to evolve quite differently.

		exposure	risk
Trend $\mu_{\mathbf{t}}$	exposure	0.00031	*−1.00000*
		(0.00002, 0.00111)	*
$\widehat{\mathbf{\Sigma}}_{\eta}$	risk	−0.00064	0.00130
		(−0.00204, −0.00002)	(0.00008, 0.00415)
Slope $\delta_{\mathbf{t}}$	exposure	0.00004	*−1.00000*
		(0.00001, 0.00009)	*
$\widehat{\mathbf{\Sigma}}_{\zeta}$	risk	−0.00007	0.00013
		(−0.00018, −0.00001)	(0.00001, 0.00039)

		policies	claims
Noise $\varepsilon_{\mathbf{t}}$	policies	0.00016	*0*
		(0, 0.00056)	*
$\widehat{\mathbf{\Sigma}}_{\varepsilon}$	claims	0	0.004206
		*	(0.00252, 0.00629)

Table 1: ML estimates of variance matrices for Victorian crash data with asymmetric 95% confidence intervals in parentheses. Correlations are given in italics on the upper triangular elements of the variance matrices.

The interventions provide therefore some flexibility in the potentially restrictive assumptions of common levels and drifts.

Figure 3 shows that the one-step ahead prediction residual series for registrations and claims are reasonably well-behaved. Normality assumptions may be considered. The residual series of claims appears to have a significant cyclical autocorrelation pattern. Extended versions of the LRT using cycles or autoregressive components can be examined to eliminate this autocorrelation. The cyclical pattern may be related to economic fluctuations in Victorian GDP or unemployment. The model may account for such effects by the inclusion of a dynamic latent

equation	year	event	estimate	t-stat
$\delta_t^{(R)}$	1970	war on 1034	−0.0785	−3.79
$\mu_t^{(R)}$	1971	seat belt law	−0.1084	−4.66
$\mu_t^{(E)}$	1980	change in data collection	−0.0864	−2.04
$\mu_t^{(R)}$	1990	advertising initiatives	−0.3757	−6.74
$\mu_t^{(E)}$	1992	change in data collection	−0.0662	−8.63

Table 2: Victorian crash data intervention estimates with t-statistics.

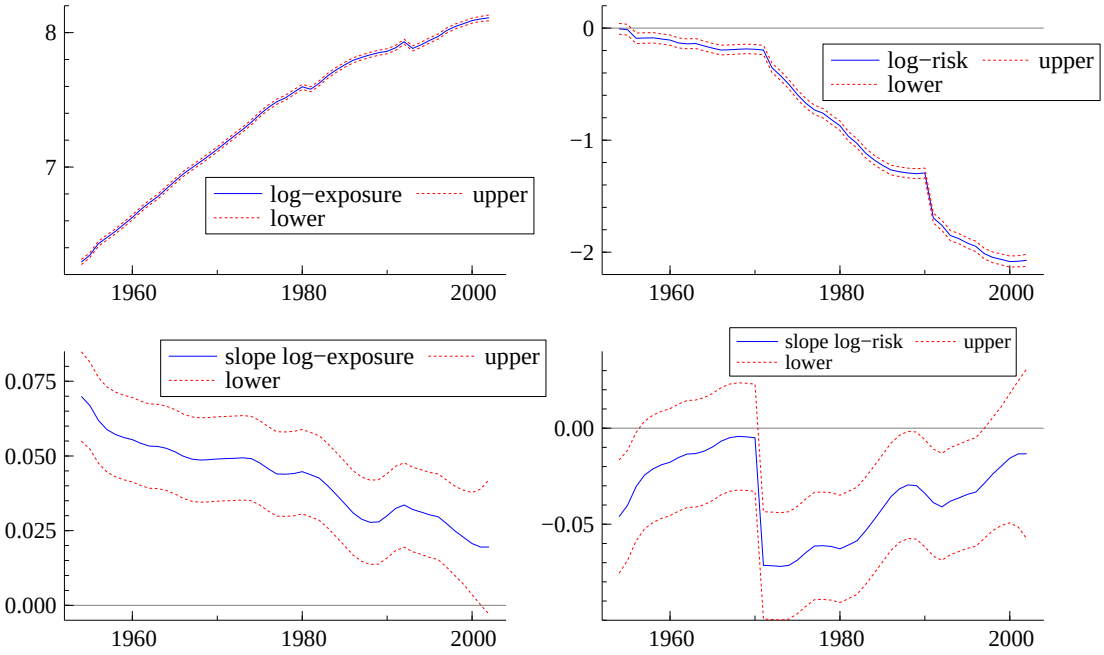

Figure 2: Smoothed output for LRT applied to Victorian insurance data

process or by explanatory variables. In this paper we concentrate on simpler structures in order to maintain the methodological focus. This application of the LRT shows how it can be effectively used for investigating insurance and road-safety issues. The model-based LRT framework can further provide point forecasts and prediction intervals for the number of claims which is useful for insurance portfolio analysis. The test procedures of Harvey (2001) can be applied to test whether a stochastic trend is required for the modelling of the risk and exposure in fatal accidents. In this analysis evidence is found to support the hypothesis of a negative relationship between risk and exposure. Changes in laws, enforcement and advertising are shown to have a significant impact on risk.

4 Case II: a three-dimensional credit card LRT model

In this section we study the usage of credit cards in Australia. The dataset consists of monthly observations for three variables: x_t, the number of credit card accounts; y_t, the number of purchases made by credit cards; z_t, the expenditure via credit cards (the total dollar value of the purchases). The observations are from May 1994 through to August 2004 (124 observations) and are presented in Figure 4. Contrary to the insurance and road crash case of the previous

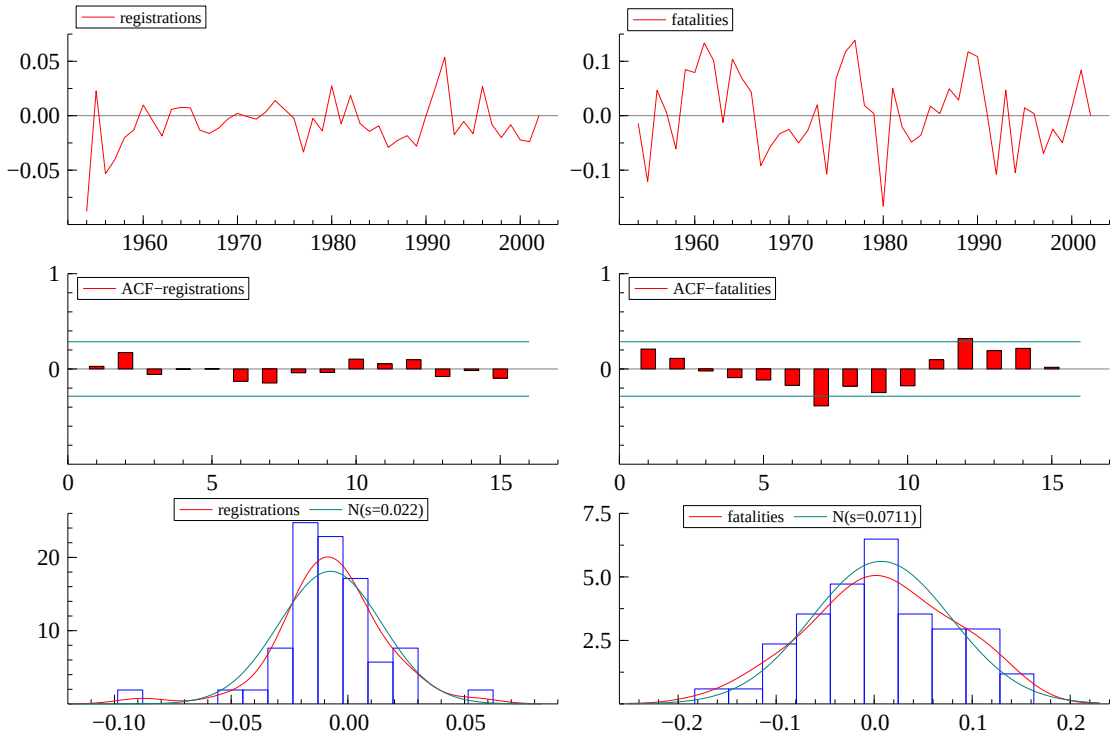

Figure 3: One-step-ahead residuals for the LRT applied to Victorian insurance data

section, we examine a system where the outcomes (the number and value of purchases) are clearly not undesirable. These events are important for Australian banks who are aggressively marketing credit cards. However, there has been recent interest in Australian consumers' reliance on credit card debt, which may be of concern to bank risk managers. The results show that the "exposure" to credit card purchases, the "risk" of purchases and the "severity" or size of purchases are increasing with a rapid growth of the total value of credit card purchases per month. The data is in nominal terms so that severity includes inflationary effects.

The application of the LRT framework can be described through the relations:

$$\begin{aligned} \text{accounts}_t &= y_t &&= E_t \times U_t^{(x)}, \\ \text{purchases}_t &= x_t &&= E_t \times R_t \times U_t^{(y)}, \\ \text{expenditure}_t &= z_t &&= E_t \times R_t \times S_t \times U_t^{(z)}, \end{aligned}$$

where E_t is exposure, R_t is risk or intensity of credit card use, S_t is severity or the value of a credit card purchase and $U_t^{(i)}$ is the multiplicative error for $i = x, y, z$. The aim is to show how the latent processes of exposure, risk and severity have developed over time and how they have influenced the total value of credit card purchases as this is the key variable for bank liquidity forecasters.

Furthermore, we examine the event of January 2002 when the Reserve Bank of Australia

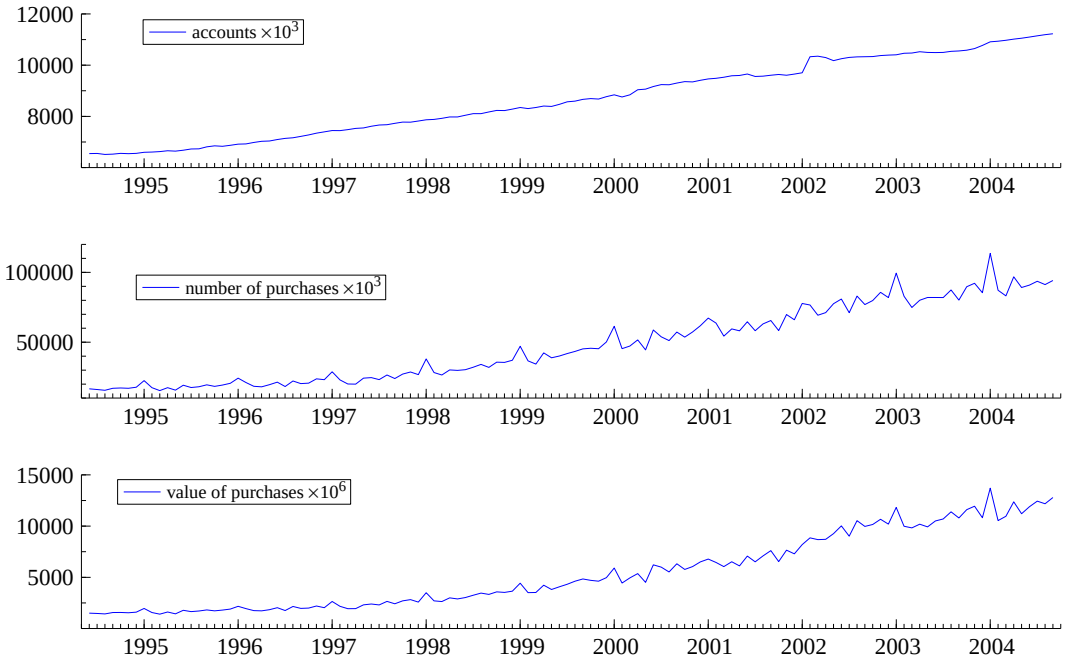

Figure 4: Credit cards observations from May 1994 through to August 2004 (monthly data): (i) x_t, number of credit card accounts; (ii) y_t, number of purchases made by credit cards; (iii) z_t, total dollar value of the purchases.

(RBA) started to include credit card accounts from commercial banks and other financial institutions in the sample. The inclusion of data from other credit card issuers means that the number of credit cards has increased but the unobserved factors risk and severity may also change since the new issuers in the sample of credit card users may represent customers with different spending patterns. The change in the composition of the sample in January 2002 is permanent and therefore level interventions are appropriate. To measure the effect of the level change in exposure, an intervention in the equation of x is included. Furthermore, level interventions are considered for the trend processes R_t and S_t. All interventions are incorporated in the LRT model as in (6).

The observations are at a monthly frequency. The seasonal variation may not affect the number of credit cards (exposure) but it may affect the observed number of purchases and expenditures since these variables are typically subject to seasonal fluctuations caused by, for example, Christmas and Easter. However we do not expect that the usage (risk) and values (severity) are affected by seasonal factors and therefore we do not integrate them in the LRT framework. In other words, we include the seasonal components in the observation equations of the y and z variables and we consider the model structure of (4) rather than (3). The LRT model specification for the credit card case with seasonal components and intervention dummies is then given by

$$\begin{array}{rcccccccccc} \log x_t & = & \mu_t^{(E)} & + & & & & & & + & \varepsilon_t^{(x)}, \\ \log y_t & = & \mu_t^{(E)} & + & \mu_t^{(R)} & + & & & \gamma_t^{(y)} & + & \varepsilon_t^{(y)} \\ \log z_t & = & \mu_t^{(E)} & + & \mu_t^{(R)} & + & \mu_t^{(S)} & + & \gamma_t^{(z)} & + & \varepsilon_t^{(z)}, \end{array} \tag{13}$$

where $\mu_t^{(E)}$ is modelled as the LLT process (2), $\mu_t^{(i)}$ $(i = R, S)$ is modelled as the LLT process with level and slope breaks (6), $\gamma_t^{(j)}$ $(j = y, z)$ is the trigonometric seasonal component and $\varepsilon_t^{(j)}$ $(j = x, y, z)$ is the observation noise term for $t = 1, \ldots, n$.

The LRT model provides a good fit of the data. The variance matrices estimates are given in Table 3. The variance matrices for trend and observation noises are treated as diagonal. This is strongly supported by the fact that the maximised loglikelihood function values have almost the same values (-934.304 for the unrestricted model and -935.495 for the restricted model). The estimated variances of the seasonal innovations are relatively large compared to the observation noise. The variances of the slope innovations are larger than the ones of the level. These estimation results lead to estimated trend components that are smooth functions of time and to estimated seasonal components that exhibit strong variations over time, see Figure 5.

All correlations in the variance matrix of the slope innovations are not significant since the value zero is part of the simulated 95% confidence bounds for the covariances. The estimate of the correlation coefficient between the slopes for exposure and severity is 0.947, indicating that

			exposure	risk	severity
Trend		exposure	1.333×10^{-5}	*0*	*0*
			$(4.888 \times 10^{-6}, 2.200 \times 10^{-5})$	*	*
	$\mathbf{\Sigma}_\eta$	risk	0	8.910×10^{-5}	*0*
			*	$(4.022 \times 10^{-7}, 3.157 \times 10^{-4})$	*
		severity	0	0	1.178×10^{-5}
			*	*	$(1.2631 \times 10^{-7}, 3.368 \times 10^{-5})$
Slope		exposure	2.605×10^{-7}	*−0.0576*	*−0.9472*
			$(2.330 \times 10^{-8}, 1.016 \times 10^{-6})$	*	*
	$\mathbf{\Sigma}_\zeta$	risk	-3.706×10^{-7}	1.587×10^{-6}	*0.3742*
			$(-1.853 \times 10^{-6}, 5.009 \times 10^{-7})$	$(4.144 \times 10^{-7}, 7.316 \times 10^{-6})$	*
		severity	-5.774×10^{-8}	5.637×10^{-8}	1.426×10^{-8}
			$(-2.604 \times 10^{-7}, 1.512 \times 10^{-7})$	$(-5.613 \times 10^{-7}, 3.885 \times 10^{-7})$	$(5.190 \times 10^{-9}, 2.997 \times 10^{-7})$
			accounts	purchases	expenditure
Seasonal		accounts	*0*	*0*	*0*
			*	*	*
	$\mathbf{\Sigma}_\xi$	purchases	0	2.202×10^{-3}	*0.9749*
			*	$(1.578 \times 10^{-3}, 3.075 \times 10^{-3})$	*
		expenditure	0	1.944×10^{-3}	1.806×10^{-3}
			*	$(1.356 \times 10^{-3}, 2.743 \times 10^{-3})$	$(1.235 \times 10^{-3}, 2.575 \times 10^{-3})$
Noise		accounts	3.125×10^{-6}	*0*	*0*
			$(4.927 \times 10^{-7}, 9.395 \times 10^{-6})$	*	*
	$\mathbf{\Sigma}_\varepsilon$	purchases	0	1.074×10^{-5}	*0*
			*	$(4.449 \times 10^{-9}, 3.793 \times 10^{-4})$	*
		expenditure	0	0	4.133×10^{-5}
			*	*	$(2.285 \times 10^{-7}, 1.658 \times 10^{-4})$

Table 3: Variance hyperparameter estimates for credit card data with 95% confidence intervals in parentheses. Correlations are given in italics on the upper triangular elements of the variance matrices.

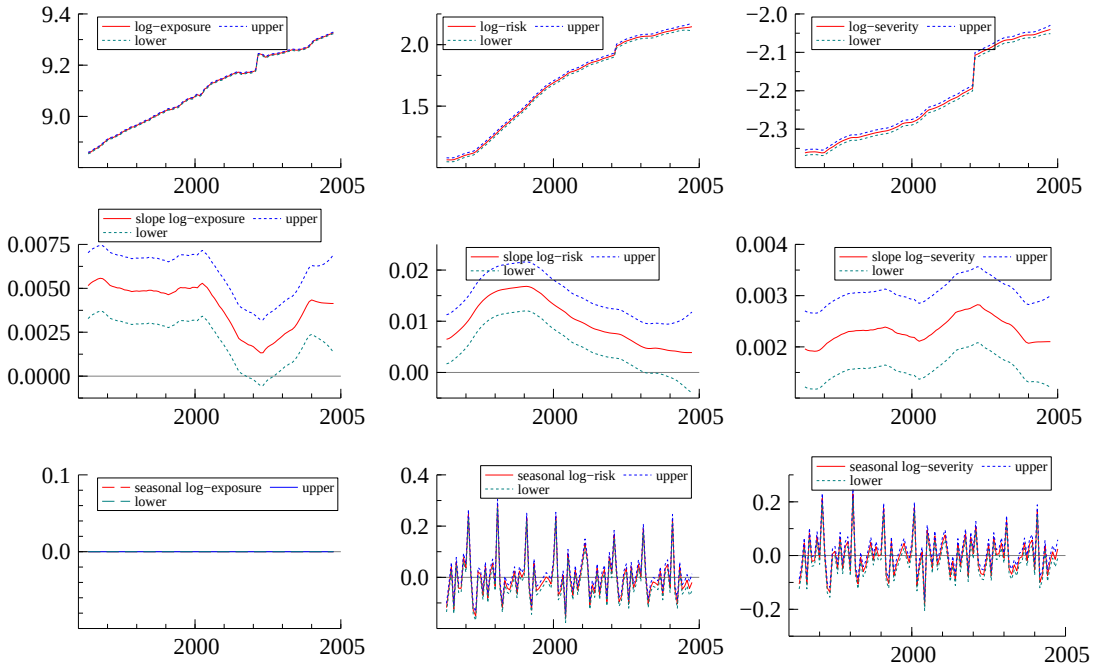

Figure 5: Smooth estimates of the latent factors in the LRT model for the Australian credit card data. The exposure, risk and severity factors are presented column-wise. The components trend, drift and seasonal are presented row-wise.

equation	month	event	estimate	t-stat
$\log x_t$	2002 Jan.	sample change	0.0615	13.05
$\log y_t$	—	—	0.0662	2.31
$\log z_t$	—	—	0.0833	8.27

Table 4: Intervention estimates for credit cards data with t-statistics

these factors move closely together. However, there is not much certainty about this feature in the set of monthly time series. The seasonal factors related to outcome (purchases) and loss (expenditure) variables are almost common since the correlation is 0.975 and its related covariance is significant with a small confidence interval. This is confirmed in Figure 5 where it is shown that the two seasonal components are very similar.

Table 4 shows that the three intervention coefficients are significant. Figure 5 illustrates the effects in the trend components. We should note that all level interventions are part of the observation equations. Although the risk factor is significantly affected by the change in survey composition, the severity of credit card purchases increased the most. It can therefore be concluded that the new account holders in the survey from January 2002 onwards are making more expensive purchases with their credit cards. The new customers have had a smaller effect on the risk (intensity) of making a purchase.

5 Case III: a multiple exposure LRT model

The yearly number of persons killed and seriously injured (KSI) in collisions between mopeds and bicycles in the Netherlands is closely watched by policy makers and the public at large since these vehicles are widely and intensively used in the Netherlands. To investigate the risk of a KSI accident, a dataset is constructed with two exposure variables and one outcome variable. The two exposure variables consist of numbers of kilometres driven by mopeds and by bicycles. No distinction is made between light mopeds and more classical mopeds. The outcome variable is the yearly number of accidents for which the primary collision partners are one moped user and one bicycle user, and for which the victims are either killed or hospitalised. The exposure variables are obtained from the Dutch national travel survey that also publish survey error variances. The latter variable is used as a precision variable of the exposure measurements. The national statistical agency supplies the outcome variable based on police records. The dataset is available for 1985–2003, at a yearly frequency (19 observations). Given this limited sample, the model needs to be parsimonious to preserve a sufficient number of degrees of freedom.

The three time series are presented in Figure 6. For the exposure series, the 95% confidence

intervals are also presented. These are based on the published survey error variances. The number of kilometres driven by bicycles are subject to a number of stepwise increases in the late 1980s and in 1994 while those by mopeds show a gradually decrease over the years. The increase in 1994 for the bicycle kilometres driven may be explained by the extension of the sample with persons under 12 years of age. The decrease of the 95% confidence intervals for the two exposure series from 1994 onwards is due to the increase of the survey sample size by a factor of two. The yearly number of accidents show stepwise decreases in 1991 and in 2000. It is anticipated that the decrease in 1991 coincides with the introduction of a free travelpass for students (typically between 17 and 21 years of age). The travelpass gave free access to the national and local public transport systems (mainly buses and trains). The usage of the free travelpass became more and more restricted over the years from 1995 onwards. This may partly explain the slow increase of KSI accidents in the late 1990s. It is reasonable to argue that the decrease in 2000 may have been caused in part by the introduction of a law that moved all mopeds from the special bicycle roads (or tracks) to the main roads in use by other motorized vehicles (motors, cars, trucks). This law only applies to situations where special bicycle roads or tracks exist and where the traffic conditions are sufficiently safe. Therefore many exceptions to this law exist and the "mopeds on the roadway" law can perhaps only partly explain this drop in 2000.

The LRT model for the number of KSI accidents caused by moped-bicycle collisions accounts for exposure and risk and is not concerned with severity. However, it remains a three-dimensional latent factor model since we have two volume or exposure variables. The structure of the LRT model is given by

$$\begin{aligned} \text{driver kilometres bicycle}_t &= E_{1t} \times U_{1t}^{(x)}, \\ \text{driver kilometres moped}_t &= E_{2t} \times U_{2t}^{(x)}, \\ \text{accidents}_t &= E_{1t} \times E_{2t} \times R_t \times U_t^{(z)}, \end{aligned} \tag{14}$$

where E_{1t} is the latent variable for bicycle exposure, E_{2t} is for moped exposure and R_t is the risk of a KSI accident. The interventions for the sample extension in 1994, to include persons under 12 years of age, affects only the bicycle exposure E_{1t} since the Dutch law forbids persons under 16 years of age to use mopeds. The introduction of the free travel pass in 1991 has a likely effect on moped exposure since the public transport system can offer an alternative to journeys made by mopeds while the bicycle is typically used for shorter travel distances. The free travel pass may also have an effect on risk and therefore a level shift in the log-risk equation at 1991 is included. Finally, the introduction of the law "mopeds on the roadway" in 2000 is incorporated in the model by a shift intervention in the log-risk equation. In summary, the

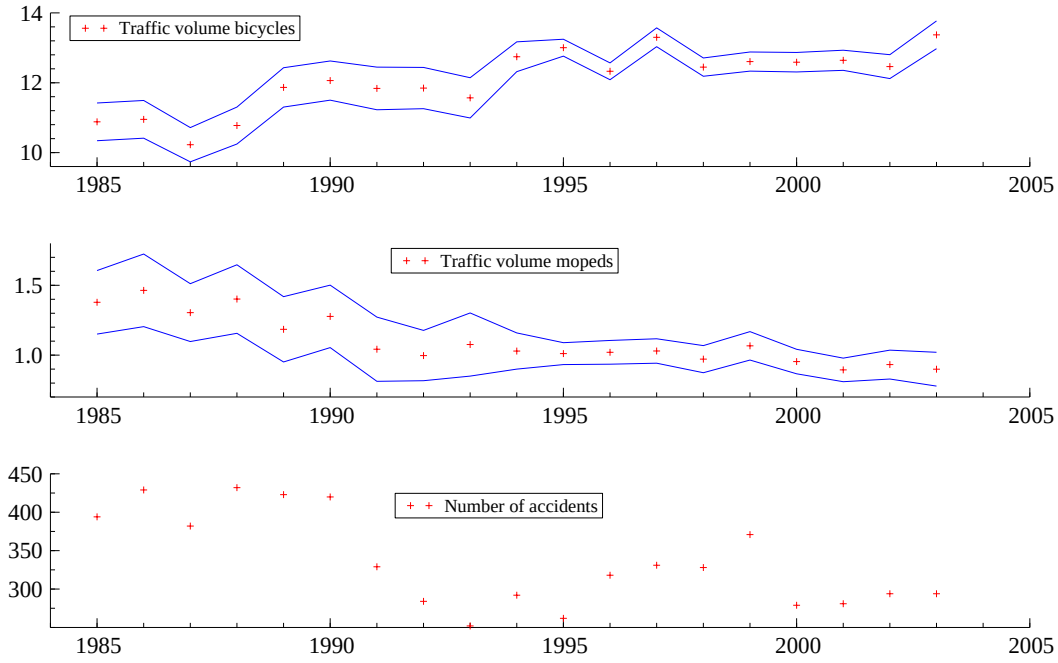

Figure 6: Moped-bicycle accidents in the Netherlands (1985 – 2003). (i) Annual traffic volumes (billion kilometres) of bicycles; (ii) Annual traffic volumes (billion kilometres) of mopeds; (iii) Annual counts of accidents (persons killed or hospitalized) between mopeds and bicycles. Note the increased accuracy o traffic volumes from 1994 onwards.

LRT model for this case is given by

$$\begin{aligned}
\log x_{1t} &= \mu_{1t}^{(E)} &&&&&&+ \varepsilon_{1t}^{x},\\
\log x_{2t} &= \mu_{2t}^{(E)} &&&&&&+ \varepsilon_{2t}^{x},\\
\log y_{t} &= \mu_{1t}^{(E)} &+ \mu_{2t}^{(E)} &+ \mu_{t}^{(R)} &+ \varepsilon_{t}^{y},
\end{aligned}$$

where the following components are subjected to shift interventions: $\mu_{1t}^{(E)}$ in 1994 (survey sample increase), $\mu_{2t}^{(E)}$ in 1991 (free travelpass), $\mu_{t}^{(R)}$ in 1991 (free travelpass) and $\mu_{t}^{(R)}$ in 2000 (law "mopeds on the roadway"). The variances of ε_{it}^{x}, for $i = 1, 2$, are set equal to a fixed non-negative parameter plus a known time-varying value that is implied by the different precisions of the survey. The variance of ε_{t}^{y} is also decomposed in this way but the time-varying value is implied by the normal approximation of the Poisson counts of accidents.

This LRT model is estimated by standard maximum likelihood methods and the estimated parameters are reported in Table 5 together with their standard errors that are computed by Monte Carlo methods. Given the short time-span of the sample, the time-variations in the level and slope components are limited. The variance matrix of the level vector, for the two exposure series and the number of accidents, is estimated as zero. This leads to estimated trend components that are smooth functions of time since the only sources of trend variations are drift changes. In the case of kilometres driven by mopeds, the slope variation is also estimated as zero and therefore we obtain a fixed time trend that is only interrupted by the estimated intervention in 1991. The covariance between the drift component of kilometres driven by bicycles and number of accidents is estimated as zero too. The estimates of the non-zero parameters are reported in Table 5. The constant variance of the observation noise for moped volume is estimated as zero. This implies that the random noise in the equation for $\log x_{2t}$ is due to the variation in the different sample sizes over the years.

Two significant intervention estimates are obtained and reported in Table 5. The introduction of the free travel pass in 1991 has a significant effect on the kilometres driven by mopeds, not on the risk factor. The law of "mopeds on the roadway" has a significant negative effect on the risk factor. The extension of the sample for bicycle volume with children under 12 years of age did not affect the analysis. We also have experimented with other possible interventions but they made little or no improvement to the likelihood function.

The smoothed estimates of the trend factors of exposure and risk are displayed in Figure 5. All estimated factors turn out to be smooth functions of time and the figures confirm the estimation results in Table 5. The interventions in the moped volume and risk factors are clearly visible and their significance is clear. The risk factor also exhibits stochastic variation. Risk is decreasing until the early 1990s, but has been increasing since 1993, as confirmed by the slope component of risk. This pattern may be explained by the popularity of light mopeds for which it is not obligatory to wear a crash helmet. It is commonly believed that many of

variances	equation	estimate	confidence interval
Drift $\boldsymbol{\Sigma}_\zeta$	exposure bicycle ($\mu_{1t}^{(E)}$)	2.73×10^{-5}	$(9.70 \times 10^{-8}, 0.00016)$
	risk ($\mu_t^{(R)}$)	0.002	$(4.94 \times 10^{-5}, 0.00709)$
Noise $\boldsymbol{\Sigma}_\varepsilon$	bicycle volume (x_{1t})	0.00097	$(0.00024, 0.00220)$
	accidents (y_t)	4.69662×10^{-5}	$(1.23 \times 10^{-5}, 0.06380)$
interventions			
1991 shift	exposure moped ($\mu_{2t}^{(E)}$)	-0.18	$(-0.29, -0.07)$
2000 shift	risk ($\mu_t^{(R)}$)	-0.31	$(-0.49, -0.13)$

Table 5: Moped versus bicycle accidents in the Netherlands (1985 – 2003). Parameter variances and interventions with 95% confidence intervals in parentheses.

the light moped vehicles are modified to enable them to drive as fast as mopeds which require a crash helmet. It is evident that accidents are likely to be more severe when the concerned moped drivers do not wear helmets. This may explain the increasing underlying trend in the number of KSI accidents.

6 Conclusions

In this paper we propose a latent risk time series (LRT) model for measuring multivariate "event risk". The model framework includes factors for exposure, risk and severity. It is general enough to have applications in fields ranging from financial risk management to road safety research. The LRT builds on existing work by providing a fully multivariate structure which allows for stochastically time-varying latent factors with possibly common dynamic features.

The multivariate nature of the model means that common state components can be identified through the state correlation structure. The magnitude and sign of correlations between states can also provide interesting interpretations for researchers. For example, the LRT model for car insurance claims indicated a strong negative relationship between accident risk and exposure, which has been supported by theory for many years.

Stochastic LRT specifications allow for time variation in parameters without requiring arbitrary re-calibration of model parameters. This is an advantage inherent in the unobserved components approach to modeling. The application to credit cards data showed that stochastic variation is important in measuring the risk and severity of credit card purchases. For the car insurance data, stochastic variation is less important — the LRT model reveals that structural breaks explain most of the changes in risk and exposure over the past 50 years.

The illustration of accidents between mopeds and bicycles has shown that the model can also include multiple categories of exposure variables. When more data is available, the LRT

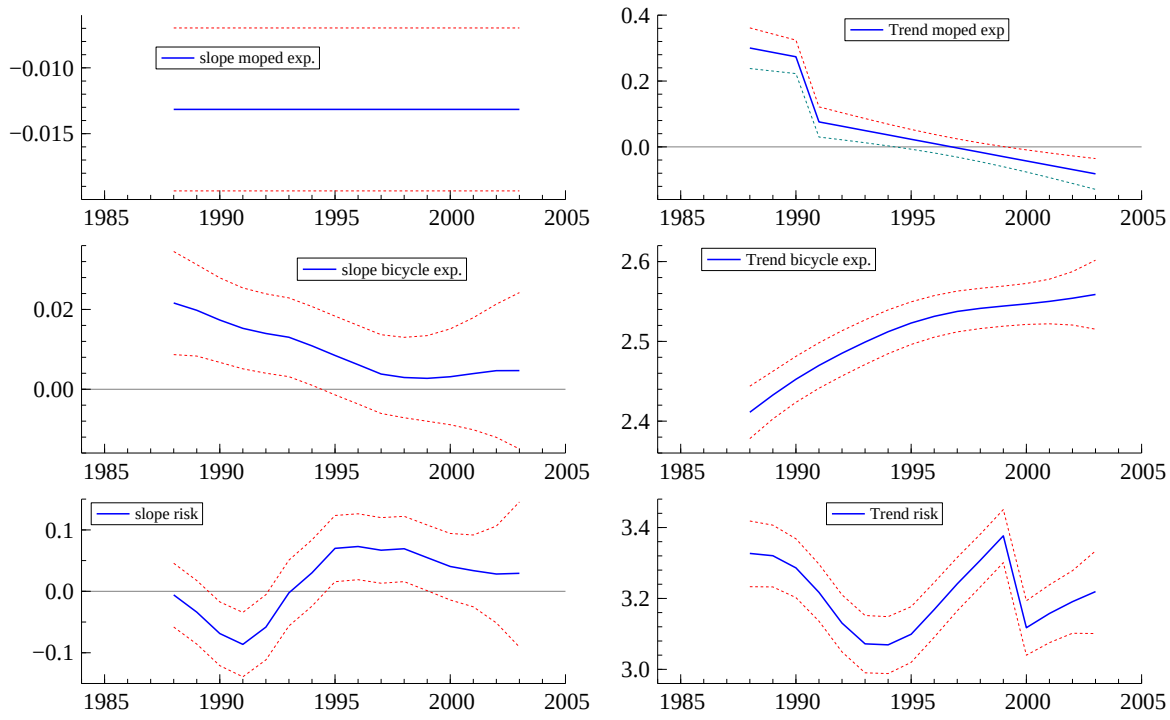

Figure 7: Smooth estimates of the latent factors in the LRT model for the moped-bicycle KSI accident data.

model can handle more detailed categories of exposure and/or risk. For example, different risk factors can be included for male/female, different age groups and different regions. Finally, a useful direction for future research is to develop methods for identifying and interpreting covariance structures in LRT models.

References

Allen, L. and A. Saunders (2003). A survey of cyclical effects in credit risk measurement models. Technical Report BIS Working Paper 126, BIS.

BIS (2004). Basel II: International convergence of capital measurement and capital standards: A revised framework. Technical Report 107, Bank of International Settlements.

de Jong, P. and P. Boyle (1983). Monitoring mortality: A state-space approach. *Journal of Econometrics 23*, 131–146.

Dominici, F., A. McDermott, and T. Hastie (2004). Improved semiparametric time series models of air pollution and mortality. *Journal of the American Statistical Association 99*, 938–948.

Durbin, J. and S. J. Koopman (2001). *Time Series Analysis by State Space Methods*. Number 24 in Oxford statistical science series. Oxford: Oxford University Press.

Fahrmeir, L. and S. Wagenpfeil (1996). Smoothing hazard functions and time-varying effects in discrete duration and competing risk models. *Journal of the American Statistical Association 91*, 1584–1594.

Finkenstadt, B. and B. Grenfell (2000). Time series modelling of childhood diseases: A dynamical systems approach. *Applied Statistics 49*, 187–205.

Gaudry, M. (1984). DRAG, un Modèle de la Demande Routière, des Accidents et leur Gravité, appliqué au Québec de 1956–1986. Technical Report Publication CRT-359, Centre de recherche sur les Transports, et Cahier #8432, Département de sciences économiques, Université de Montréal.

Gaudry, M. and S. Lassarre (Eds.) (2000). *Structural Road Accident Models: The International DRAG Family*. Oxford: Elsevier Science Ltd.

Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter*. Cambridge: Cambridge University Press.

Harvey, A. C. (2001). Testing in unobserved components models. *Journal of Forecasting 20*, 1–19.

Harvey, A. C. and J. Durbin (1986). The effects of seat belt legislation on British road casualties: A case study in structural time series modelling. *Journal of the Royal Statistical Society A 149*(3), 187–227.

Harvey, A. C. and C. Fernandes (1989). Time series models for insurance claims. *Journal of the Institute of Actuaries 116*, 513–528.

Ledolter, J., S. Klugman, and C. Lee (1990). Credibility models with time-varying trend components. *ASTIN Bulletin 21*.

Li, L. and K. Kim (2000). Estimating driver crash risk based on the extended bradley-terry model: An induced exposure method. *Journal of the Royal Statistical Society A 163*, 227–240.

Morton, A. and B. Finkenstadt (2005). Discrete time modelling of disease incidence time series by using markov chain monte carlo methods. *Applied Statistics 54*, 575–594.

Oppe, S. (1989). Macroscopic models for traffic and traffic safety. *Accident Analysis and Prevention 21*, 225–232.

WHO (2004). World report on road traffic injury prevention. Edited by M. Peden, R. Scureld, D. Sleet, D. Mohan, A. A. Hyder, E. Jarawan and C. Mathers.